



(11)

EP 3 528 125 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention
of the grant of the patent:
07.09.2022 Bulletin 2022/36

(51) International Patent Classification (IPC):
G06F 11/30^(2006.01) G06F 11/07^(2006.01)
G06F 11/20^(2006.01)

(21) Application number: **18211717.6**

(52) Cooperative Patent Classification (CPC):
G06F 11/0751; G06F 11/0793; G06F 11/2015;
G06F 11/3003

(22) Date of filing: **11.12.2018**

(54) **POWER SUPPLY UNIT FAN RECOVERY PROCESS**

LÜFTERWIEDERHERSTELLUNGSPROZESS EINER STROMVERSORGUNGSEINHEIT

PROCÉDÉ DE RÉCUPÉRATION DE VENTILATEUR D'UNITÉ D'ALIMENTATION ÉLECTRIQUE

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB
GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO
PL PT RO RS SE SI SK SM TR

(30) Priority: **30.01.2018 US 201862623694 P**
27.07.2018 US 201816047382

(43) Date of publication of application:
21.08.2019 Bulletin 2019/34

(73) Proprietor: **Quanta Computer Inc.**
Taoyuan City 33377 (TW)

(72) Inventors:
• **CHEN, Chao-Jung**
33377 Taoyuan City (TW)

- **CHEN, Yi-Chieh**
33377 Taoyuan City (TW)
- **WU, Yueh-Chang**
33377 Taoyuan City (TW)
- **SHIH, Ching-Yi**
33377 Taoyuan City (TW)

(74) Representative: **dompatent von Kreisler Selting**
Werner -
Partnerschaft von Patent- und Rechtsanwälten
mbB
Deichmannhaus am Dom
Bahnhofsvorplatz 1
50667 Köln (DE)

(56) References cited:
US-A1- 2002 101 714 US-A1- 2011 227 521
US-A1- 2012 069 514

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 3 528 125 B1

Description

FIELD

[0001] The present invention relates to controlling operation of a power supply unit, and more specifically, to systems and methods for adjusting operation of a power supply unit which is not meeting the needs of a computer system.

BACKGROUND

[0002] Computer systems typically include a large number of computer components. The components need to be working properly in order to ensure the successful, continuous operation of the computer system. Power supply units (PSUs) have unique roles in computer systems because PSUs provide power and also include an attached cooling fan to circulate air over heat-producing components in the computer system. Although the PSU's cooling fan is typically weaker than other fans in the computer system, but the PSU's cooling fan can still provide crucial cooling. For example, in some computer systems, the failure of a PSU or the PSU's fan has been known to cause damage due to overheating. Therefore, it is extremely important to the functionality of a computing system that a PSU is quickly restored to operation (particularly its fan) if it fails to function properly.

[0003] The placement of a PSU can affect the ability to restore functionality of the PSU. For example, in a first typical configuration, the PSU and its fan can be placed along the exterior edge of the computer system while stronger system fans are placed in front of or behind the PSU fan. In second typical configuration, the PSU and its fan can be placed in parallel with the stronger system fans. The second configuration can be problematic for PSU fans because the stronger system fans can cause a reverse air current through the weaker PSU fans. For example, if the PSU fan is operating at less than full capacity or fails for any reason, air can be pulled between the blades of the PSU fan and pulled through the input of the stronger fan. This can cause the weaker PSU fan to slow down, cease operating, fail to restart, and even rotate in the incorrect direction.

[0004] Although the first configuration can be implemented in many computer systems, other computer systems can require the second configuration due to space, shape, sizing, and wiring restrictions for the computer components in such computer systems.

[0005] In view of the foregoing, systems and methods are needed to detect when a PSU in a computer system is not operating properly and to enable the PSU to return to normal operation, i.e., so that that a cooling fan in the PSU can return to having an adequate operating speed and providing a correct airflow direction. In US 2002/101714 A1, a computer system fan is described. According to this document, a sensor can be provided to detect a fan obstruction, wherein a fan monitoring unit

might attempt to clear the obstruction by altering the rotation of the fan. In addition, in US 2011/227521 A1, a fan starting method is described. According to the described method, a fan module includes a control unit and a fan, wherein the control unit determines whether the fan rotates within a predetermined time. If the fan does not rotate, an alarm signal or an off signal is issued by the control unit.

10 SUMMARY

[0006] In order to solve the above-mentioned problem, a computer-implemented method for controlling two or more power supply units of a server system as defined in independent claim 1 is suggested. Further preferred embodiments of the present invention are defined in the dependent claims. The various examples of the present disclosure are directed towards a computer-implemented method for controlling two or more PSUs of a server system. An exemplary method, according to an embodiment of the present disclosure, runs through a series of steps to check whether a PSU is functioning properly, and then restores proper functionality according to the system's requirements if the PSU is found to be not functioning properly. In one embodiment, the method comprises first determining whether the two or more PSUs are switched on to an AC power source. If the two or more PSUs are found not switched on to an AC power source, the method further comprises restarting the AC power source for the two or more PSUs and completing a software-based recovery process. If the PSUs are found to be switched on to an AC power source, the method further comprises determining whether the two or more PSUs meet a predefined criterion. If the two or more PSUs do meet a predefined criterion, the method further comprises disabling a cold-redundancy mode for the two or more PSUs. If the two or more PSUs do not meet a predefined criterion, the method further comprises starting a wake-up process for a first PSU from a cold-redundancy mode.

[0007] In one exemplary implementation, the predefined criterion can include an uncoupled state of the two or more PSUs, and at least one of the following: failure of at least one fan in the server system, and/or a temperature in the computer system that is greater than a threshold operational temperature.

[0008] In another exemplary implementation, the wake-up process can comprise uncoupling the two or more PSUs and evaluating each PSU individually to identify whether each PSU is operating properly. Determining if each PSU is operating properly can comprise identifying when both of the following conditions are met: (1) a power input to the first PSU is greater than a first amount and (2) a power output of a first PSU is equal to a second amount and then checking if the first PSU has failed. If the PSU has failed, the method can further comprise completing a software-based recovery process for the first PSU. If the first PSU has not failed, then the method

can further comprise proceeding to a normal operating mode for the two or more PSUs. The first amount can be twenty-five watts. The second amount can be zero watts.

[0009] In another exemplary implementation, the software-based recovery process can comprise reducing the duty of at least one system fan to a predetermined percentage for a predetermined period of time. The method can then set a checkpoint period where the two or more PSUs cannot send a failure alert. The method can then comprise starting a power-capping function wherein power is reduced for at least one other power component of the server system. The length of the checkpoint period can be ten seconds. The predetermined percentage can be forty percent. The predetermined period of time can be ten seconds.

[0010] The words "computer system," "computing system," and "server system" are all used interchangeably in this present disclosure, and can be used to identify any electronic computing system that can use a fan to cool overheating electronic components. Such an electronic computing system can include, but not be limited to, a personal computer, a laptop computer, a tablet, and a commercial or private server system.

[0011] The words "fan" and "cooling fan" are used interchangeably in this present disclosure, and can be used to identify an electrically powered machine creating a flow of air in a preferred direction.

[0012] Reducing a fan's "duty" refers to reducing a selected speed for the fan.

[0013] Additionally, PSU fans are always attached to a corresponding PSU when described in the present disclosure.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014]

FIG. 1A shows a schematic diagram of a computing system where PSU fans are placed behind a system fan.

FIG. 1B shows a schematic diagram of a computing system where PSU fans are placed in front of a system fan.

FIG. 2 shows a schematic diagram of a computing system where PSU fans are placed beside and in parallel with a system fan.

FIG. 3A is a flowchart of the methodology for controlling two or more PSUs of a server system.

FIG. 3B is a flowchart of the methodology for an exemplary software-based recovery process according to an embodiment of the present disclosure.

FIG. 3C is a flowchart of the methodology for an exemplary wake-up process according to an embodiment of the present disclosure.

FIGs. 4 and 5 illustrate exemplary systems, in accordance with various examples of the present disclosure.

DETAILED DESCRIPTION

[0015] The present invention is described with reference to the attached figures, wherein like reference numerals are used throughout the figures to designate similar or equivalent elements. The figures are not drawn to scale and they are provided merely to illustrate the instant invention. Several aspects of the invention are described below with reference to example applications for illustration. It should be understood that numerous specific details, relationships, and methods are set forth to provide a full understanding of the invention.

[0016] The present disclosure is directed to a methodology that continually monitors a computer system to enable restoration of an incorrectly operating PSU, particularly its fan. The system first checks to identify the status of the PSUs in the system. More than one PSU can be coupled together in a cold-redundancy mode to provide back-up in case one of the PSU fails. When restoring the failed PSU, the methodology also needs to take into account potential reverse airflow from stronger system fans.

[0017] Therefore, an exemplary method, according to an embodiment of the present disclosure, will first reduce the duty of the system fans to a predetermined percentage for a predetermined period of time. This reduction will solve the problem of reverse airflow into the PSU fan because the system fans will no longer generate a current so strong that it prevents the PSU fan from starting. For example, the duty of the system fans can be reduced to forty percent for ten seconds. This reduction has been tested in computer systems and shown to be an adequate reduction.

[0018] The method will additionally set a black-out window to prevent the PSU from sending out a failure notification during the black-out window. The black-out window can be a predetermined length of time during which a check is run at the beginning and end of the period to see whether the PSU has begun functioning properly again. For example, the predetermined length of time can be ten seconds. For example, an initial warning can be sent out if the PSU fan does not reach a target speed within 5 seconds. A second warning can be sent out at 15 seconds to indicate that the PSU fan has still not reached its target speed.

[0019] This method successfully restores proper functionality to a PSU through its continual monitoring of the computer system; ability to solve the problem of a reverse airflow; and reliance on accurate percentages for system fan speed reduction and power-capping for other system components. This method provides a unique solution to a failed or failing PSU through its use of a protection process which is software-based, and does not require additional hardware components. Furthermore, the suggested percentages and time periods for adjustment provide a new and efficient protocol for restoring operation to the PSU while having a minimum impact on other system components.

[0020] FIG. 1A shows a first exemplary arrangement

for fans in a computer system. This arrangement shows one possible placement for PSU fans where multiple PSU fans can be found in line behind at least one system fan. FIG. 1A shows two PSU fans, but there can be any number of PSU fans, so long as there is at least one. In the arrangement of FIG. 1A, no reverse airflow occurs after a PSU failure because the larger system fan is always drawing air in the same direction as the PSU fans are designed to blow. Therefore, any air current generated from the system fan will be in the same direction as the intended direction for air current through PSU fan's. Exemplary air current paths are displaced by the arrows running through the components.

[0021] FIG. 1B shows a second exemplary arrangement for fans in a computer system which would also not cause reverse airflow. In the arrangement of FIG. 1B, the PSU is placed directly in front of the system fan. In the arrangement of FIG. 1B, no reverse airflow occurs after a PSU failure because the larger system fan is always drawing air in the same direction as the PSU fans are designed to blow. Therefore, any air current generated from the system fan will be in the same direction as the intended direction for air current through PSU fan's.

[0022] In contrast to the arrangement of FIGs. 1A and 1B, FIG. 2 shows an exemplary arrangement in a computer system where the placement of the PSU fans and the system fans can lead to a reverse airflow that can impede the performance of the PSU. In some cases, such an arrangement can even prevent the PSU from returning to proper operation. Exemplary air current paths are displaced by the arrows running through the components. In FIG. 2, arrows 202 running through the PSU fans demonstrate the proper way that air should flow through the PSU fan. Similarly, arrows 204 running through the system fans demonstrate the proper way that air should flow through the system fans.

[0023] However, when the system fans are configured as shown in FIG. 2, the system fans can pull air in from a surrounding region behind the system fan. Such system fans can pull air in from multiple directions (as indicated by hatched regions 206). As a result, if an area of a system fan is sufficiently close to an air inlet for a PSU, the system fan can pull air through the PSU (as indicated by arrows 208). Thus, a reverse air current (shown by arrows 210) can form through the PSU fan unit in response to the strong air draw from the system fan. Consequently, FIG. 2 shows the need for a solution to prevent the PSU fan from rotating in the reverse direction when PSU fans and system fans are in this configuration. In addition, if the PSU fan ever stops rotating, the PSU fan might not be able to restart operation because the reverse airflow current created by the system can be too strong for the weaker PSU fan to overcome.

[0024] FIG. 3A is an exemplary flow chart of steps for implementing a method 300A according to the present disclosure which can restart operation of a PSU that is not operating properly. In step 302, method 300A begins with a baseboard management controller (BMC) moni-

toring the functionality of its PSUs. In step 304, a determination is made as to whether two or more PSUs are switched onto an AC power source.

[0025] If the two or more PSUs are switched onto an AC power source, the BMC then proceeds to step 310 to check whether the PSUs meet a predefined cold-redundancy criterion. An exemplary cold-redundancy criterion can examine whether the PSUs are in an uncoupled state and whether there is failure of at least one fan in the server system, and/or a temperature in the computer system that is greater than a threshold operational temperature.

[0026] If the predefined criterion is met at step 310, then the BMC proceeds to step 312 where the cold-redundancy mode is disabled. The BMC can then return to step 302 to continue monitoring the system. The BMC will run through the methodology 300A again to ensure that the PSU is operating properly. If the predefined criterion is not met at step 310, then method 300A proceeds to method 300C where the computer system begins a wake-up process from the cold-redundancy mode. Method 300C is discussed below in more detail with respect to FIG. 3C.

[0027] This wake-up process uncouples the PSUs and proceeds to evaluate each PSU individually to see which one is not operating properly. This wake-up process is shown in FIG. 3C and begins in step 332. At step 332, a first determination is made as to whether a power input to a first PSU is less than a first amount, and whether a power output of the first PSU is equal to a second amount. If answer to both components of the first determination is "NO", then method 300C proceeds back to step 304 in method 300A of FIG. 3A. This indicates that the PSUs are operating properly and the method then continues to monitor the PSU system.

[0028] If the first determination is true, the method 300C can proceed to a second determination which evaluates whether a power input is greater than a first amount and the power output is greater than a second amount. If answer to both components of the second determination is "NO", then method 300C proceeds back to step 304 in method 300A of FIG. 3A. This indicates that the PSUs are operating properly and the method then continues to monitor the PSU system.

[0029] However, if answer to both of these is "YES", then method 300C proceeds to step 336. In step 336, a determination is made as to whether a PSU has failed. When a PSU fails, the PSU can release an alert to the computer system to identify that the PSU has failed. An alert allows the computer system to detect a failure before the computer system overheats. If, in step 336, it is determined that a PSU has failed, method 300C proceeds method 300B, particularly step 322, in order to complete the software-based recovery process for the failed PSU. Method 300B is presented below in greater detail with respect to FIG. 3B. In contrast, if it is determined at step 336 that there are no failed PSUs, method 300C proceeds to step 340. At step 340, the method checks that

the PSU has been switched from a coupled cold-redundancy mode to an uncoupled, normal operation mode. The method can then stop here or return to the beginning of method 300A in FIG. 3A to continue monitoring the computer system.

[0030] As noted above, method 300C will proceed to step 322 in method 300B of FIG. 3B to perform a software-based recovery process if a failed PSU status is found in step 336. Steps 322, and 324 define the software-based recovery process of FIG. 3B. In the software-based recovery process of method 300B, the PSU is power-cycled to restart it. Because of the reverse airflow discussed earlier, power-cycled PSUs do not always successfully restart because the PSU fan can be too weak to spin in the intended direction. Therefore, the software-based recovery process shown in FIG. 3B can increase the chance that the PSU will restart. This process is described in further detail below.

[0031] First, at step 322 of the software-based recovery process, the duty of any system fans will be reduced to a predetermined percentage for a predetermined period of time. Reducing the duty of a system fan reduces the speed with which the system fan rotates, and will thus reduce any problematic reverse airflow current. In some exemplary methods, one can reduce the duty to forty percent for ten seconds. Any other percentage and period of time can be used as well, so long as the combination of the percentage and period of time comprises a sufficient reduction to reduce a reverse airflow current in the computer system. The reduction of the reverse airflow current can be sufficient to allow the PSU fan to begin rotating in the intended direction.

[0032] After the reduction in duty at step 322, method 300B proceeds to step 324 of the software-based recovery process. In step 324, a "black-out window" is set which prevents the PSU from sending a "fault" command. For example, the "black-out window" will be a period of time during which the PSU cannot send a failure alert. This allows the method to proceed with restoring the PSU to normal operation without receiving regular alerts that the PSU is not functioning properly yet. For example, the black-out period can be ten seconds long. The black-out period also does not have to start instantly when the system fan duty is reduced in step 322. For example, step 324 can occur after a five second delay from step 322.

[0033] The software-based recovery process of 300B can be initiated under other conditions. That is, without requiring execution of method 300C. For example, referring back to FIG. 3A, and step 304 in particular, if two or more PSUs are found not switched onto an AC power source, method 300A will proceed to step 306. In step 306, the computer system will evaluate whether the AC power source for the PSU needs to be power-cycled. If it does not need to be power-cycled, the method then proceeds to the beginning of the method 300A to continue monitoring the status of PSUs in the computer system.

[0034] If the AC power source for the PSU does need to be power-cycled, method 300A provides for power-

cycling the PSU (step 308) and then proceeding to the software-based recovery process of method 300B. The software-based recovery process was previously discussed in detail with regards to steps 322, and 324.

[0035] All steps of the methods of FIGs. 3A-3C can be completed by a processor configured to complete the steps in a computer system.

[0036] A brief introductory description of example systems and networks, as illustrated in FIGs. 4-5, is disclosed herein. These variations shall be described herein as the various examples are set forth. The present disclosure now turns to FIG. 4.

[0037] FIG. 4 illustrates an example computing system 400, in which components of the computing system are in electrical communication with each other using a bus 402. The system 400 includes a processing unit (CPU or processor) 430, and a system bus 402 that couples various system components, including the system memory 404 (e.g., read only memory (ROM) 406 and random access memory (RAM) 408), to the processor 430. The system 400 can include a cache of high-speed memory connected directly with, in close proximity to, or integrated as part of the processor 430. The system 400 can copy data from the memory 404 and/or the storage device 412 to the cache 428 for quick access by the processor 430. In this way, the cache can provide a performance boost for processor 430 while waiting for data. These and other modules can control or be configured to control the processor 430 to perform various actions. Other system memory 404 may be available for use as well. The memory 404 can include multiple different types of memory with different performance characteristics. The processor 430 can include any general purpose processor and a hardware module or software module, such as module 1 414, module 2 416, and module 3 418 embedded in storage device 412. The hardware module or software module is configured to control the processor 430, as well as a special-purpose processor where software instructions are incorporated into the actual processor design. The processor 430 may essentially be a completely self-contained computing system, and containing multiple cores or processors, a bus, memory controller, cache, etc. A multi-core processor may be symmetric or asymmetric.

[0038] To enable user interaction with the computing device 400, an input device 420 is provided as an input mechanism. The input device 420 can comprise a microphone for speech, a touch-sensitive screen for gesture or graphical input, keyboard, mouse, motion input, and so forth. In some instances, multimodal systems can enable a user to provide multiple types of input to communicate with the system 400. In this example, an output device 422 is also provided. The communications interface 424 can govern and manage the user input and system output.

[0039] Storage device 412 can be a non-volatile memory to store data that are accessible by a computer. The storage device 412 can be magnetic cassettes, flash

memory cards, solid state memory devices, digital versatile disks, cartridges, random access memories (RAMs) 408, read only memory (ROM) 406, and hybrids thereof.

[0040] The controller 410 can be a specialized micro-controller or processor on the system 400, such as a BMC. In some cases, the controller 410 can be part of an Intelligent Platform Management Interface (IPMI). Moreover, in some cases, the controller 410 can be embedded on a motherboard or main circuit board of the system 400. The controller 410 can manage the interface between system management software and platform hardware. The controller 410 can also communicate with various system devices and components (internal and/or external), such as controllers or peripheral components, as further described below.

[0041] The controller 410 can generate specific responses to notifications, alerts, and/or events, and communicate with remote devices or components (e.g., electronic mail message, network message, etc.) to generate an instruction or command for automatic hardware recovery procedures, etc. An administrator can also remotely communicate with the controller 410 to initiate or conduct specific hardware recovery procedures or operations, as further described below.

[0042] The controller 410 can also include a system event log controller and/or storage for managing and maintaining events, alerts, and notifications received by the controller 410. For example, the controller 410 or a system event log controller can receive alerts or notifications from one or more devices and components, and maintain the alerts or notifications in a system event log storage component.

[0043] Flash memory 432 can be an electronic non-volatile computer storage medium or chip that can be used by the system 400 for storage and/or data transfer. The flash memory 432 can be electrically erased and/or reprogrammed. Flash memory 432 can include EPROM (erasable programmable read-only memory), EEPROM (electrically erasable programmable read-only memory), ROM, NVRAM, or CMOS (complementary metal-oxide semiconductor), for example. The flash memory 432 can store the firmware 434 executed by the system 400, when the system 400 is first powered on, along with a set of configurations specified for the firmware 434. The flash memory 432 can also store configurations used by the firmware 434.

[0044] The firmware 434 can include a Basic Input / Output System or equivalents, such as an EFI (Extensible Firmware Interface) or UEFI (Unified Extensible Firmware Interface). The firmware 434 can be loaded and executed as a sequence program each time the system 400 is started. The firmware 434 can recognize, initialize, and test hardware present in the system 400 based on the set of configurations. The firmware 434 can perform a self-test, such as a POST (Power-on-Self-Test), on the system 400. This self-test can test functionality of various hardware components such as hard disk

drives, optical reading devices, cooling devices, memory modules, expansion cards, and the like. The firmware 434 can address and allocate an area in the memory 404, ROM 406, RAM 408, and/or storage device 412, to store an operating system (OS). The firmware 434 can load a boot loader and/or OS, and give control of the system 400 to the OS.

[0045] The firmware 434 of the system 400 can include a firmware configuration that defines how the firmware 434 controls various hardware components in the system 400. The firmware configuration can determine the order in which the various hardware components in the system 400 are started. The firmware 434 can provide an interface, such as an UEFI, that allows a variety of different parameters to be set, which can be different from parameters in a firmware default configuration. For example, a user (e.g., an administrator) can use the firmware 434 to specify clock and bus speeds; define what peripherals are attached to the system 400; set monitoring of health (e.g., fan speeds and CPU temperature limits); and/or provide a variety of other parameters that affect overall performance and power usage of the system 400. While firmware 434 is illustrated as being stored in the flash memory 432, one of ordinary skill in the art will readily recognize that the firmware 434 can be stored in other memory components, such as memory 404 or ROM 406.

[0046] System 400 can include one or more sensors 426. The one or more sensors 426 can include, for example, one or more temperature sensors, thermal sensors, oxygen sensors, chemical sensors, noise sensors, heat sensors, current sensors, voltage detectors, air flow sensors, flow sensors, infrared thermometers, heat flux sensors, thermometers, pyrometers, etc. The one or more sensors 426 can communicate with the processor, cache 428, flash memory 432, communications interface 424, memory 404, ROM 406, RAM 408, controller 410, storage device 412 and display 436, via the bus 402, for example. The one or more sensors 426 can also communicate with other components in the system via one or more different means, such as inter-integrated circuit (I2C), general purpose output (GPO), and the like. Different types of sensors (e.g., sensors 426) on the system 400 can also report to the controller 410 on parameters, such as cooling fan speeds, power status, operating system (OS) status, hardware status, and so forth.

[0047] FIG. 5 illustrates an example computer system 500 having a chipset architecture that can be used in executing the described method(s) or operations, and in generating and displaying a graphical user interface (GUI). Computer system 500 can include computer hardware, software, and firmware that can be used to implement the disclosed technology. System 500 can include a processor 510, representative of a variety of physically and/or logically distinct resources capable of executing software, firmware, and hardware configured to perform identified computations. Processor 510 can communicate with a chipset 502 that can control input to and output from processor 510. In this example, chipset 502 outputs

information to output device 514, such as a display, and can read and write information to storage device 516, which can include magnetic media, and solid state media, for example. Chipset 502 can also read data from and write data to RAM 518. A bridge 504 for interfacing with a variety of user interface components 506, can be provided for interfacing with chipset 502. Such user interface components 506 can include a keyboard, a microphone, touch detection and processing circuitry, a pointing device, such as a mouse, and so on. In general, inputs to system 500 can come from any of a variety of sources, machine generated and/or human generated.

[0048] Chipset 502 can also interface with one or more communication interfaces 508 that can have different physical interfaces. Such communication interfaces can include interfaces for wired and wireless local area networks, for broadband wireless networks, and for personal area networks. Further, the machine can receive inputs from a user via user interface components 506 and execute appropriate functions, such as browsing functions by interpreting these inputs using processor 510.

[0049] Moreover, chipset 502 can also communicate with firmware 512, which can be executed by the computer system 500 when powering on. The firmware 512 can recognize, initialize, and test hardware present in the computer system 500 based on a set of firmware configurations. The firmware 512 can perform a self-test, such as a POST, on the system 500. The self-test can test the functionality of the various hardware components 502-518. The firmware 512 can address and allocate an area in the memory 518 to store an OS. The firmware 512 can load a boot loader and/or OS, and give control of the system 500 to the OS. In some cases, the firmware 512 can communicate with the hardware components 502-510 and 514-518. Here, the firmware 512 can communicate with the hardware components 502-510 and 514-518 through the chipset 502, and/or through one or more other components. In some cases, the firmware 512 can communicate directly with the hardware components 502-510 and 514-518.

[0050] It can be appreciated that example systems 400 and 500 can have more than one processor (e.g., 430, 510), or be part of a group or cluster of computing devices networked together to provide greater processing capability.

[0051] For clarity of explanation, in some instances, the present disclosure may be presented as including individual functional blocks including functional blocks, including devices, device components, steps or routines in a method embodied in software, or combinations of hardware and software.

[0052] In some embodiments, the computer-readable storage devices, mediums, and memories can include a cable or wireless signal containing a bit stream and the like. However, when mentioned, non-transitory computer-readable storage media expressly exclude media such as energy, carrier signals, electromagnetic waves, and signals per se.

[0053] Methods according to the above-described examples can be implemented using computer-executable instructions that are stored or otherwise available from computer readable media. Such instructions can include, for example, instructions and data which cause or otherwise configure a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions. Portions of computer resources used, can be accessible over a network. The computer executable instructions may be, for example, binaries and intermediate format instructions, such as assembly language, firmware, or source code.

[0054] Devices implementing methods according to these disclosures can include hardware, firmware and/or software, and can take any of a variety of form factors. Typical examples of such form factors include laptops, smart phones, small form factor personal computers, personal digital assistants, rack-mount devices, standalone devices, and so on. Functionality described herein also can be embodied in peripherals or add-in cards. Such functionality can also be implemented on a circuit board among different chips, or different processes executing in a single device, by way of further example.

[0055] While various examples of the present invention have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of the present invention should not be limited by any of the above described examples. Rather, the scope of the invention should be defined in accordance with the following claims and their equivalents.

[0056] Although the invention has been illustrated and described with respect to one or more implementations, equivalent alterations and modifications will occur to others skilled in the art upon the reading and understanding of this specification and the annexed drawings.

[0057] The terminology used herein is for the purpose of describing particular examples only and is not intended to be limiting of the invention. As used herein, the singular forms "a," "an," and "the" are intended to include the plural forms as well, unless the context clearly indicates otherwise. Furthermore, to the extent that the terms "including," "includes," "having," "has," "with," or variants thereof, are used in either the detailed description and/or the claims, such terms are intended to be inclusive in a manner similar to the term "comprising."

[0058] Unless otherwise defined, all terms (including technical and scientific terms) used herein have the same meaning as commonly understood by one of ordinary skill in the art to which this invention belongs. It will be further understood that terms, such as those defined in commonly used dictionaries, should be interpreted as having a meaning that is consistent with their meaning in the context of the relevant art and will not be interpreted in an idealized or overly formal sense unless expressly so defined herein.

Claims

1. A computer-implemented method (300A) for controlling two or more power supply units (PSUs) of a server system, each PSU comprising an internal fan, and the server system comprises a system fan, wherein the computer-implemented method comprising:

monitoring the PSUs by a baseboard management controller (BMC) and detecting a failure of a PSU, said detecting being performed by checking whether the PSUs are switched on to an AC power source;

if the PSUs are switched on to the AC power source, then a predefined criterion is checked to determine if the PSUs are in an uncoupled state and at least one of the following is true: a failure of at least one fan in the server system and/or a temperature in the server system is greater than a threshold operational temperature;

if the PSUs are not switched on to the AC power source, restarting the AC power source for the two or more PSUs and performing a software-based recovery process which is comprised of:

reducing duty of at least one system fan to a predetermined percentage for a predetermined period of time (322); and
setting a checkpoint period where the two or more PSUs cannot send a failure alert.

2. The computer-implemented method of claim 1, wherein if the PSUs are switched onto the AC power source, and the predefined criterion is not met, the method further comprises:

starting a wake-up process from a cold-redundancy mode said wake-up process comprising uncoupling the PSUs and evaluating the proper operation of each PSU individually, wherein the two or more PSUs are coupled together in the cold-redundancy mode.

3. The computer-implemented method of claim 2, wherein starting the wake-up process for the first PSU from the cold-redundancy mode further comprises:

checking a power change status, wherein the power change status can include either:

when a power input to the first PSU is less than a first amount and when a power output of the first PSU is equal to a second amount; or

when a power input to the first PSU is greater than a first amount and when a power output of the first PSU is greater than a second amount.

4. The computer-implemented method of claim 3,

wherein starting the wake-up process for the first PSU from the cold-redundancy mode further comprises determining that the first PSU has failed.

5. The computer-implemented method of claim 3, wherein starting the wake-up process for the first PSU from the cold-redundancy mode further comprises:

determining that the first PSU has not failed; detecting that the software-based recovery process is not needed; and
in response to detecting that the software-based recovery process is not needed, proceeding to a normal operating mode for the two or more PSUs.

6. The computer-implemented method of claim 5, wherein the first amount is twenty-five watts.

7. The computer-implemented method of claim 5, wherein the second amount is zero watts.

8. The computer-implemented method of any of claims 1 to 7, wherein if the PSUs are switched on to the AC power source, and the predefined criterion is met, the method further comprises:

disabling a cold-redundancy mode of the two or more PSUs, wherein the two or more PSUs are coupled together in the cold-redundancy mode.

9. The computer-implemented method of any of claims 1 to 8, wherein the predetermined period of time is ten seconds.

10. The computer-implemented method of any of claims 1 to 9, wherein the checkpoint period is ten seconds.

11. The computer-implemented method of any of claims 1 to 10, wherein the predetermined percentage is forty percent.

Patentansprüche

1. Computerimplementiertes Verfahren (300A) zum Steuern zweier oder mehr Netzteileinheiten (PSUs) eines Serversystems, wobei jede PSU ein inneres Gebläse aufweist, und wobei das Serversystem ein Systemgebläse aufweist, wobei das computerimplementierte Verfahren die folgenden Schritte aufweist:

Überwachen der PSUs durch einen Baseboard-Management-Controller (BMC) und Erkennen eines Ausfalls einer PSU, wobei das Erkennen durchgeführt wird, indem geprüft wird, ob die PSUs auf eine Wechselstromquelle geschaltet

sind;

wenn die PSUs auf die Wechselstromquelle geschaltet sind, wird ein vordefiniertes Kriterium geprüft, um festzustellen, ob die PSUs sich in einem nicht gekoppelten Zustand befinden und mindestens eine der folgenden Bedingungen zutrifft: ein Ausfall mindestens eines Gebläses in dem Serversystem und/oder eine Temperatur in dem Serversystem ist höher als eine Schwellen-Betriebstemperatur;

wenn die PSUs nicht auf die Wechselstromquelle geschaltet sind, Neustarten der Wechselstromquelle für die zwei oder mehr PSUs und Durchführen eines softwarebasierten Wiederherstellungsprozesses, welcher umfasst:

Verringern der Auslastung mindestens eines Systemgebläses auf einen vorbestimmten Prozentsatz für einen vorbestimmten Zeitraum (322); und

Einstellen eines Kontrollpunktzeitraums, in welchem die zwei oder mehr PSUs keinen Ausfallalarm senden können.

2. Computerimplementiertes Verfahren nach Anspruch 1, bei welchem, wenn die PSUs auf die Wechselstromquelle geschaltet sind und das vordefinierte Kriterium nicht erfüllt ist, wobei das Verfahren ferner den folgenden Schritt aufweist:

Starten eines Prozesses des Aufweckens aus einem kalten Redundanzmodus, wobei dieser Aufwachprozess das Entkoppeln der PSUs und das individuelle Auswerten des ordnungsgemäßen Betriebs jeder PSU aufweist,

wobei die zwei oder mehr PSUs in dem kalten Redundanzmodus miteinander gekoppelt sind.

3. Computerimplementiertes Verfahren nach Anspruch 2, bei welchem das Starten des Prozesses des Aufweckens der ersten PSU aus dem kalten Redundanzmodus ferner die folgenden Schritte aufweist:

Prüfen eines Stromänderungszustands, wobei der Stromänderungszustand sein kann, entweder:

wenn ein Stromeingang in die erste PSU geringer als eine erste Menge ist und wenn ein Stromausgang der ersten PSU gleich einer zweiten Menge ist; oder

wenn ein Stromeingang in die erste PSU größer als eine erste Menge ist und wenn ein Stromausgang der ersten PSU größer als eine zweite Menge ist.

4. Computerimplementiertes Verfahren nach Anspruch 3, bei welchem das Starten des Vorgangs des Aufweckens der ersten PSU aus dem kalten Redundanzmodus ferner das Feststellen des Ausfalls der ersten PSU aufweist.

undanzmodus ferner das Feststellen des Ausfalls der ersten PSU aufweist.

5. Computerimplementiertes Verfahren nach Anspruch 3, bei welchem das Starten des Vorgangs des Aufweckens der ersten PSU aus dem kalten Redundanzmodus ferner die folgenden Schritte aufweist:

Feststellen, dass die erste PSU ausgefallen ist; Erkennen, dass der softwarebasierte Wiederherstellungsprozess nicht erforderlich ist; und in Reaktion auf die Feststellung, dass der softwarebasierte Wiederherstellungsprozess nicht erforderlich ist, Fortschreiten zu einem Normalbetriebsmodus für die zwei oder mehr PSUs.

6. Computerimplementiertes Verfahren nach Anspruch 5, bei welchem die erste Menge fünfundzwanzig Watt beträgt.

7. Computerimplementiertes Verfahren nach Anspruch 5, bei welchem die zweite Menge null Watt beträgt.

8. Computerimplementiertes Verfahren nach einem der Ansprüche 1 bis 7, bei welchem, wenn die PSUs auf die Wechselstromquelle geschaltet sind und das vordefinierte Kriterium erfüllt ist, das Verfahren den folgenden Schritt aufweist:

Deaktivieren eines kalten Redundanzmodus der zwei oder mehr PSUs, wobei die zwei oder mehr PSUs im kalten Redundanzmodus miteinander gekoppelt sind.

9. Computerimplementiertes Verfahren nach einem der Ansprüche 1 bis 8, bei welchem der vorbestimmte Zeitraum zehn Sekunden beträgt.

10. Computerimplementiertes Verfahren nach einem der Ansprüche 1 bis 9, bei welchem der Kontrollpunktzeitraum zehn Sekunden beträgt.

11. Computerimplementiertes Verfahren nach einem der Ansprüche 1 bis 10, bei welchem der vorbestimmte Prozentsatz vierzig Prozent beträgt.

Revendications

1. Procédé mis en œuvre par ordinateur (300A) pour contrôler deux blocs d'alimentation (PSU) ou plus d'un système serveur, chaque PSU comprenant un ventilateur interne et le système serveur comprend un ventilateur système, dans lequel le procédé mis en œuvre par ordinateur comprend :

la surveillance des PSU par un contrôleur de

- gestion de carte de base (BMC) et la détection d'une panne d'un PSU, ladite détection étant effectuée en vérifiant si les PSU sont commutés sur une source d'alimentation CA ;
- si les PSU sont commutés sur la source d'alimentation CA, un critère prédéfini est ainsi vérifié pour déterminer si les PSU sont dans un état découplé et si au moins l'une des conditions suivantes est vraie : une panne d'au moins un ventilateur dans le système serveur et/ou une température dans le système serveur est supérieure à une température de fonctionnement seuil ;
- si les PSU ne sont pas commutés sur la source d'alimentation CA, le redémarrage de la source d'alimentation CA pour les deux ou plusieurs PSU et le fait d'effectuer un processus de récupération basé sur logiciel qui est constitué de :
- la réduction du fonctionnement d'au moins un ventilateur du système à un pourcentage prédéterminé pendant une période de temps prédéterminée (322) ; et
 - la définition d'une période de point de vérification, les deux PSU ou plus ne pouvant pas envoyer d'alerte de panne.
2. Procédé mis en œuvre par ordinateur selon la revendication 1, dans lequel si les PSU sont commutés sur la source d'alimentation CA et le critère prédéfini n'est pas satisfait, le procédé comprend en outre :
- le démarrage d'un processus de réveil à partir d'un mode de redondance froide, ledit processus de réveil comprenant le découplage des PSU et l'évaluation du bon fonctionnement de chaque PSU individuellement, les deux PSU ou plus étant couplés ensemble dans le mode de redondance froide.
3. Procédé mis en œuvre par ordinateur selon la revendication 2, dans lequel le démarrage du processus de réveil pour le premier PSU à partir du mode de redondance froide comprend en outre : la vérification d'un état de changement de puissance, l'état de changement de puissance pouvant comporter soit :
- lorsqu'une entrée de puissance vers le premier PSU est inférieure à une première quantité et lorsqu'une puissance de sortie du premier PSU est égale à une seconde quantité ; ou
 - lorsqu'une entrée de puissance vers le premier PSU est supérieure à une première quantité et lorsqu'une puissance de sortie du premier PSU est supérieure à une seconde quantité.
4. Procédé mis en œuvre par ordinateur selon la revendication 3, dans lequel le démarrage du processus de réveil pour le premier PSU à partir du mode de redondance froide comprend en outre la détermination que le premier PSU n'est pas défaillant ;
- la détection que le processus de récupération basé sur le logiciel n'est pas nécessaire ; et
- en réponse à la détection que le processus de récupération basé sur le logiciel n'est pas nécessaire, la procédure à un mode de fonctionnement normal pour les deux PSU ou plus.
5. Procédé mis en œuvre par ordinateur selon la revendication 3, dans lequel le démarrage du processus de réveil pour le premier PSU à partir du mode de redondance froide comprend en outre :
- la réduction du fonctionnement d'au moins un ventilateur du système à un pourcentage prédéterminé pendant une période de temps prédéterminée (322) ; et
 - la définition d'une période de point de vérification, les deux PSU ou plus ne pouvant pas envoyer d'alerte de panne.
6. Procédé mis en œuvre par ordinateur selon la revendication 5, dans lequel la première quantité est de vingt-cinq watts.
7. Procédé mis en œuvre par ordinateur selon la revendication 5, dans lequel la seconde quantité est de zéro watt.
8. Procédé mis en œuvre par ordinateur selon l'une quelconque des revendications 1 à 7, dans lequel si les PSU sont commutés sur la source d'alimentation CA et que le critère prédéfini est satisfait, le procédé comprend en outre :
- la désactivation d'un mode de redondance froide des deux PSU ou plus, les deux PSU ou plus étant couplés ensemble dans le mode de redondance froide.
9. Procédé mis en œuvre par ordinateur selon l'une quelconque des revendications 1 à 8, dans lequel la période de temps prédéterminée est de dix secondes.
10. Procédé mis en œuvre par ordinateur selon l'une quelconque des revendications 1 à 9, dans lequel la période de point de vérification est de dix secondes.
11. Procédé mis en œuvre par ordinateur selon l'une quelconque des revendications 1 à 10, dans lequel le pourcentage prédéterminé est de quarante pour cent.




FIG. 1A




FIG. 1B




FIG. 2




FIG. 3A




FIG. 3B




FIG. 3C




FIG. 4




FIG. 5

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 2002101714 A1 [0005]
- US 2011227521 A1 [0005]